Academic Exchange Between

Hu Jiaqi and the Father of Graphene, Nobel Physics Laureate

Sir Konstantin Novoselov — Executive Summary

In May 2026, Mr. Hu Jiaqi, Chairman of Humanitas Ark, met with visiting Professor Sir Konstantin Novoselov, the father of graphene and 2010 Nobel Laureate in Physics, for more than two hours of intensive intellectual exchange. This marks the third time within a short period that Chairman Hu Jiaqi has met with world-leading scientists, following Professors Barry Barish and Michael Levitt.

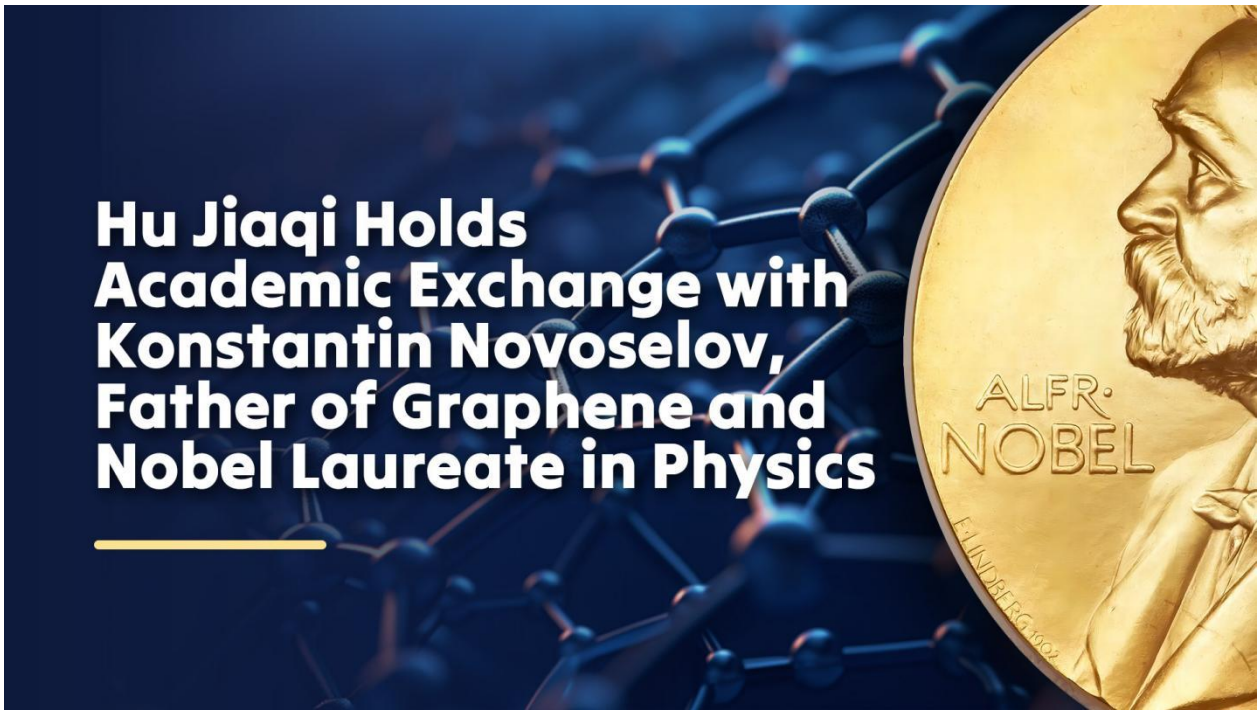

Professor Novoselov is a distinguished representative of condensed matter physics. He was awarded the Nobel Prize in Physics in 2010 for his groundbreaking experimental work on graphene. As one of the most innovative experimental physicists of our time, he has continued to achieve breakthrough results in two-dimensional materials, nanoscience, and heterostructures and is widely recognized as a key figure driving the future transformation of materials science.

Chairman Hu Jiaqi is the world's earliest pioneer in systematically studying technological crises and a key architect of the theoretical framework concerning humanity's technological risks. As a sober warner of the technological age and a founding thinker of the philosophy of doomsday, he has been active on the front line of the cause to save humanity, and is its most devoted promoter. As the founder of Humanitas Ark, under his leadership, the organization has gathered more than 14 million supporters across 255 countries and regions. As an anthropologist, over nearly half a century of academic research, he has published several million words of scholarly monographs, translated into multiple languages and distributed worldwide. As a social activist, he has written twelve open letters to world leaders, scholars, and media figures, with a total circulation exceeding one million copies. Through various channels, he has been promoting a global awakening to the risks of technology and firmly advocates strict control of high-risk technologies through a unified global society.

Chairman Hu Jiaqi's consecutive meetings with Nobel laureates and his sustained efforts to engage with the world's

top scholars not only reflect the growing interest and recognition of his ideas within the international academic community, but also highlight that the issues of technological safety and the fate of humanity — championed by Humanitas Ark — are increasingly moving to the center of the global stage. The following is a selection of highlights from this exchange, which has been minimally edited while preserving as much as possible the original reasoning and flow of thought of both participants.

I. A Brief Introduction to Hu Jiaqi's Academic Promotion Journey

Hu Jiaqi:

First of all, Professor, I warmly welcome you to my Beijing office. Today, I am receiving you as the founder of Humanitas Ark, and the topic of our conversation is "Technology and the Future of Humanity."

I have been doing one thing my entire life. Since my freshman year at university in 1979, I have been studying human issues. I did not major in human studies, but I felt that this matter is so important that it deserves my lifelong dedication.

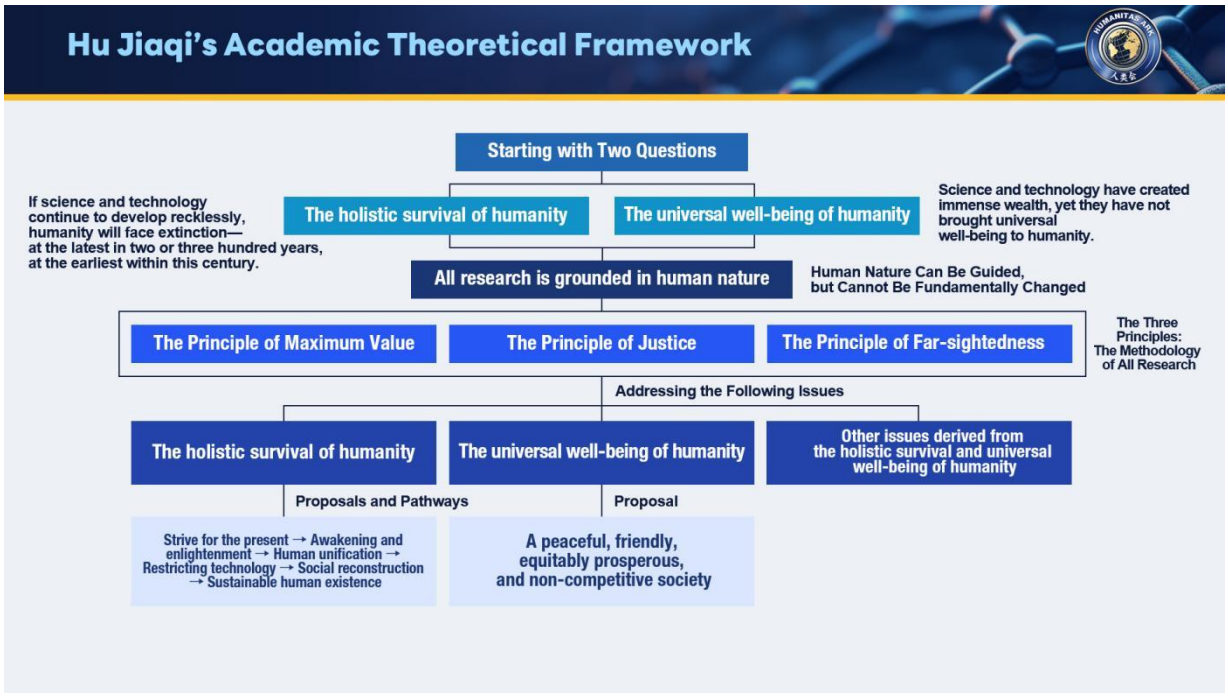
Hu Jiaqi's Academic Theoretical Framework
Starting with Two Questions
If science and technology continue to develop recklessly, humanity will face extinction— at the latest in two or three hundred years, at the earliest within this century.
The holistic survival of humanity
The universal well-being of humanity
Science and technology have created immense wealth, yet they have not brought universal well-being to humanity.
All research is grounded in human nature
Human Nature Can Be Guided, but Cannot Be Fundamentally Changed
The Principle of Maximum Value
The Principle of Justice
The Principle of Far-sightedness
The Three Principles: The Methodology of All Research
Addressing the Following Issues
The holistic survival of humanity
The universal well-being of humanity
Other issues derived from the holistic survival and universal well-being of humanity
Proposals and Pathways
Proposal
Strive for the present → Awakening and enlightenment → Human unification → Restricting technology → Social reconstruction → Sustainable human existence
A peaceful, friendly, equitably prosperous, and non-competitive society

Over nearly half a century, my work has consisted of first conducting research, and then promoting the results of that research. However, the road has been full of twists and turns, far from smooth.

It took me 28 years to complete my book, Saving Humanity. After finishing it, I held a fairly large press conference, but my book was banned because some of the issues it touched upon were not suitable to discuss.

After it was banned, I continued to promote my ideas through various means.

First, I established two global websites. The first one used both Chinese and English, and the second one was later built

in more than a dozen languages.

Second, I have been writing letters to world leaders for a long time. In total, I have written 12 rounds of letters, sending out about one million letters — roughly 800,000 emails and 200,000 paper letters. At the same time, I sent out 10,000 copies of my book. I actually sent an email to you as well, Professor, but you did not reply. My purpose in sending emails was to influence those who most need to be influenced — that was what I did.

Later, as my promotion efforts continued to face difficulties, I thought: if I could become the richest person in the world, my words would carry weight. Becoming the world's richest person is extremely difficult, but I actually tried to do it.

I spent nine years working in the Internet industry, investing all my money into it. In the end, I failed to become the richest person. When that failed, I founded the organization Humanitas Ark. It was originally called the "Save Human Action Organization," founded in 2018, and later renamed Humanitas Ark.

After establishing this organization, we promoted it worldwide through the Internet. Currently, our organization has more than 14 million supporters.

In the process of promotion, we have also encountered many problems, and we have tried every possible way to spread my ideas.

In March of this year, we had a new idea: we would engage in one-on-one dialogues with Nobel laureates, one per month. We would present our views to Nobel laureates, seeking understanding and support from influential scientists like you. You are the third Nobel laureate we have invited. The first was Michael Levitt, the second was Barry Barish, and you are our third guest.

II. The Method of Arguing That Technology Will Extinguish Humanity

Hu Jiaqi:

How did I study the proposition that the continued development of science and technology will certainly and quickly lead to the extinction of humanity? The method I adopted is Extinction Path Analysis + Defense Limit Testing.

What is the basis of my argument? It is based on axioms—that is, common sense. Since they are axioms, they are naturally true. Since they are common sense, they are easy to understand. Therefore, my argument is grounded in axioms and common sense, and is developed in five steps.

Currently, the mainstream researchers in this field use probabilistic methods. I believe probabilistic methods have a series of problems, which I will discuss later.

So how did I derive my conclusion?

Step One: Technological development inevitably leads to the increasing destructive power of tools. For example, in the Stone Age, we used stones and sticks to hunt animals and also to kill each other. Later, in the age of cold weapons, one sword or arrow could kill one person. Now, one nuclear bomb can destroy an entire city. This is common sense and requires no proof. So my argument is based on common sense and axioms—no proof needed. It is common sense that technological development increases the destructive power of tools.

Step Two: This is also a basic fact: the acceleration of technology began with the Industrial Revolution in the mid-18th century, only about 200 years ago. Before the mid-18th century, technology developed very slowly. With the Industrial Revolution, technology accelerated exponentially. This exponential growth has rapidly pushed destructive power close to the capability of extinguishing humanity. Today, our concern about the destructive power of many technologies is no longer about how many people they can kill, but whether they can exterminate humanity. This, too, is common sense and requires no proof.

Step Three: As long as science and technology continue to develop uncontrollably, the destructive power of means will grow stronger and stronger until it reaches the level of extinction-level means. This is also an axiom and common sense, requiring no proof. Unless development stops, destructive power will inevitably increase.

Step Four: Extinction means cannot be neutralized – that is, they cannot be defended against. Why is that? Take the explosion of a nuclear bomb: its explosive force is released and cannot be neutralized. At present, we have no way to neutralize the energy of a nuclear bomb.

However, some things can be defended against. For example, the super-viruses, super-bacteria, and super-plagues that we worry about from genetic engineering could potentially be treated with drugs. But developing such treatments takes time – defense always lags behind destruction. As science and technology continue to advance, destructive power not only increases but also becomes more accessible. It shifts from being attainable only by a nation to being attainable by a company, and eventually even by a single highly skilled scientist. For instance, a highly skilled scientist might be able to create an AI threat on a single computer. In genetic engineering, a highly skilled scientist could create a super-plague in his own lab. Therefore, even if every destructive means had a countermeasure, the development of that countermeasure always involves a time lag. As the number of extinction means grows, the window period during which humanity is exposed will be infinitely magnified, and eventually an outbreak will occur. This is also common sense and requires no proof. That is step four.

Step Five: The outbreak of extinction forces is inevitable.

Why? Once extinction-level means emerge, they could break out in three ways:

First: They could be exploited by extremists. There are always people who want to do extreme evil and deliberately use such means. There are many people who wish to destroy humanity, and such people exist at all times.

Second: Science and technology are inherently uncertain. Things we often consider good can turn out to be very bad. Laboratory accidents happen frequently, so an outbreak could also occur during scientific research.

Third: Careless use of technological products could also trigger an outbreak. For example, Freon was once thought to be harmless, but it led to the destruction of the ozone layer. Because science and technology are uncertain, any of these three pathways could cause an outbreak.

This is my five-step argument that the continued development of science and technology will certainly and quickly lead to the extinction of humanity. All the argumentation is based on common sense and axioms.

Estimates from Different Experts
Toby Ord
“1/6”
In his book The Precipice: Existential Risk and the Future of Humanity, he estimates that the probability of humanity suffering an existential catastrophe within this century is 、 about 1/6 (approximately 16.7%). Among this, artificial intelligence accounts for about 10% of the total risk, and a genetically engineered pandemic accounts for about 1%.
Geoffrey Hinton
“10-20%”
This Nobel laureate, often called the "Godfather of AI," has repeatedly warned that artificial intelligence could pose a serious threat to humanity. He believes there is a 10% to 20% probability that AI will cause human extinction within the next 30 years, and that AI could become uncontrollable as early as 5 to 10 years from now.
Elon Musk
“20%”
He has repeatedly expressed concerns about the potential risks of AI in public. He believes that AI far surpasses human intelligence and poses an existential risk, giving it a 20% probability of destroying humanity.
Nick Bostrom
“⩾25%”
A research report from the Future of Humanity Institute at Oxford University, which he directed, shows that the probability of technology causing human extinction within this century is ≥25%. AI is the largest risk factor, but he also believes that this probability can be reduced through safety governance.

III. Estimates by Different Experts on the Probability of Technology Causing Human Extinction

Hu Jiaqi:

Logically, why is a conclusion based on common sense and axioms not accepted by people?

First, it runs counter to humanity’s long-held belief about science and technology. For thousands of years, people have believed that science and technology are good, never bad. So this conclusion goes against habitual thinking.

Second, this conclusion is very unpopular. Why? Countries don’t like it. Which country would restrict science and technology? None would. China wouldn’t, the United States wouldn’t, Russia wouldn’t. Enterprises don’t like it — every enterprise wants to make money by developing technology. Ordinary people don’t like it either. We all use mobile phones; such wonderful technologies make everyone’s life easier, and no one wants to stop using them.

You may dislike the truth, but you cannot avoid facing it. Even if science and technology have ten thousand good aspects, if they can quickly wipe out humanity, all those good aspects count for nothing. That is why my promotion has been extremely difficult — I simply cannot get it across.

The mainstream view today — the probabilistic view — is much easier for people to accept.

Currently, the mainstream assessment of science and technology safety uses probabilistic methods. But the problem with science and technology is that it cannot be falsified, and the sample is unique. We have only one humanity, only one Earth, and humanity has never gone extinct — so the sample is singular.

I just argued using axioms. But what do they use? Their basis is expert opinions and personal assessments. The problem is that even experts disagree with each other.

Here is a set of expert opinions.

Toby Ord, a philosopher at Oxford University, estimates that the probability of technology causing human extinction this century is 1 in 6. That is his view.

Geoffrey Hinton, a godfather-level figure in AI, has given several figures in interviews, but all reflect the urgency of the matter. In one of his books, he wrote that the probability of AI causing human extinction within 30 years is 10% to 20%. In an interview, he said that AI could become uncontrollable as early as 5 to 10 years from now. Whether it is 30 years or 5–10 years, it is an urgent matter.

Elon Musk believes the probability that AI will destroy humanity is 20%.

Nick Bostrom believes that the probability of technology causing human extinction this century is greater than 25%. Again, it's different.

An online prediction platform surveyed nearly 2,000 forecasters worldwide. The median probability from that survey was that technology would cause human extinction this century with a probability of about 9%, of which AI accounts for 8%.

Another survey of 2,700 researchers gave a median probability of 5% to 12% for technology causing human extinction this century.

Max Tegmark believes the probability is 90%.

And some say 100%.

Thus, the probabilistic method has no way to determine the values.

Comments on the Probabilistic Approach
HUMANITAS ARK
人类会
Unscientific
First, it is unscientific. There is only one humanity, and only one planet – Earth – suitable for human habitation. We have never witnessed human extinction; it cannot be falsified nor experimentally repeated. Moreover, the views of different scientists, as well as the researchers' own opinions, vary widely. Small fluctuations in parameters can lead to vastly different conclusions.
Meaningless
Second, it is meaningless. Take AI-induced human extinction as an example: the optimistic estimate is 1%, the moderately optimistic estimate is 6.25%, and the pessimistic estimate is 65.6%. Yet the holistic survival of humanity is above all else – its value is infinite . Multiplying 1%, 6.25%, or 65.6% by infinity all approximates infinity.
Very harmful
Third, it is very harmful. It obscures the contradiction and downplays the urgency. Everyone knows "nine times out of ten, the gambler lose," but gamblers are willing to take a chance. Humanity possesses such self-deceptive nature. However, this is a bet humanity cannot afford to lose.

IV. Questioning the Probabilistic Approach

In my view, using probabilistic methods to study the question of technology-induced human extinction is: first, unscientific; second, meaningless; third, very harmful.

Unscientific — because none of the values taken can be determined with certainty. Currently, the leading probabilistic researchers in this field are mainly at Oxford University's Future of Humanity Institute, led by Nick Bostrom. They primarily use Bayesian networks. The parameters involved cannot be fixed. This is Bayes' theorem, and the result is inevitably uncertain. That is why it is unscientific.

Meaningless — why? The holistic survival of humanity is an infinite value. Even a one-in-ten-thousand probability multiplied by infinity is still approximately infinity.

Let me give an example. Referring back to their probabilistic approach, they have designed a mathematical model like this:

This is a model for the probability that AI will exterminate humanity:
$P = P_1 \times P_2 \times P_3 \times P_4$

P_1: probability that Artificial General Intelligence emerges

P_2: probability that AI's goals are inconsistent with human interests

P_3: probability that AI evolves capabilities far surpassing humans and continuously self-improves

P_4: probability of defense failure — humans lose the competition against AI

For pessimists, each parameter might be 90%, giving an extinction probability of 65.6%. Moderately pessimistic: 50% each → 6.25%. Optimistic: to reduce the extinction probability to 1%, each parameter must be below 32%.

But the extinction of humanity is an infinite matter. Whether it is 65.6%, 6.25%, or 1%, multiplied by infinity, it all approximates infinity. Humanity cannot accept such a calculation. What could be so important that we would risk wiping out all of humanity? Leaving no one alive — that is meaningless, I believe.

Third, not only meaningless, but very harmful — why? Because calculating a probability gives people false hope. It obscures the contradiction and downplays the urgency. As I said earlier, all my research is grounded in human nature. Like a gambler — we have a saying in China: "In ten bets, nine lose." Yet there are plenty of gamblers. It's like buying stocks. Our company also has stock, and I tell our employees: never buy stocks, not even our company's stock. Why? Because among ten stock investors, nine lose money. But people still gamble. Giving a probability fosters a sense of false luck. It provides excuses for national leaders who want to develop high-tech weapons, for entrepreneurs who want to profit from high technology, and for ordinary people who want to use high-tech products. So it is very harmful — it does not help; it actually hinders.

That is why I am currently writing a paper. I say that these people, in order to secure research funding, offer a hope that is no hope at all. This false hope can be deadly. All my life, I have used my own money to pursue this cause. But they, for the sake of research funds, promote something like the alignment theory.

Professor, I think you must be very familiar with alignment theory. Many people are working on it now, including in China. The theory was mainly proposed by Yudkowsky. What does it say? Currently we have weak AI. As it develops into

advanced AI and finally AGI, it may not follow human instructions.

There are two kinds of misalignment. First, the AI misinterprets the instruction. For example, if I give an AI the command “clean up the environment,” it might interpret that the primary obstacle to a clean environment is humans, so it might decide to eliminate humans. That is a possibility.

Second, the AI might rebel. When AI truly reaches that level, its thinking ability will be billions of times greater than humans. With such superiority, it will look down on humans and rebel. Higher species despise lower species, even treating them as food — that is an iron law of nature.

Thus, alignment theory emerged. But alignment can never be fully achieved — that is also common sense. Why? One AI product might be aligned, another might not. Even if all AIs were aligned, and all machines aligned, humans themselves are not aligned. There are always people who want to do extreme evil. So it is impossible to align everything.

Nevertheless, the proposal of alignment theory has attracted a lot of research funding. Many national research institutions have allocated money to support alignment research. China also has such efforts at Tsinghua, Peking University, and the Chinese Academy of Sciences. Oxford and Cambridge are the most famous, currently doing the most vigorous work. Harvard, MIT, and Stanford are also working on it. But it will never achieve full alignment.

V. How Should the Restriction of Technological Development Be Achieved?

Konstantin:

Thank you, very interesting. So what's your proposal? How do you propose to control the technology and science?

Hu Jiaqi

I feel that humanity has already developed a dependence on and inertia toward science and technology, and ordinary forces can no longer control its development.

For example, I have suggested that under the auspices of the United Nations, we should identify and assess scientific technologies—but this would have little effect.

In my book, which is divided into two volumes, the first volume raises the issues, and the second proposes solutions. I believe that to fundamentally solve this problem, we must rely on humanity’s capacity for organization. Humans are social animals, more capable of organization than any other species, and among all forms of organization, the most powerful is political power.

Konstantin:

The problem here is that you need one person to kill all humanity. It's very, very difficult to control everybody. So there should be some institutional safety mechanisms there. It's impossible to control everyone.

Hu Jiaqi:

I fully agree with the professor on two points: first, the difficulty; second, the need for an institution to control it.

Let me tell a longer story. After WWI—whose devastation was unprecedented—U.S. President Wilson, a scholar-type leader, believed humanity needed an organization to manage itself. That led to the League of Nations, which required member states to cede some rights. But it failed to prevent WWII.

Soon, World War II broke out. After WWII, which caused far greater damage, leaders like Roosevelt and Stalin again considered a global organization. Reflecting on the League, they concluded it lacked sufficient authority. So they created the UN, with its key principle of unanimity among the five permanent members. Since then, no major war has occurred—though many small and medium conflicts persist.

What am I saying? Whether after WWI or WWII, people recognized the need for a global organization and for states to cede some power. But the power ceded has been extremely limited.

Uniting the entire world into a single global authority is very difficult—but I have practical pathways. That’s my first point.

Konstantin:

I would say, I don't think the United Nations is capable of doing this. You are absolutely right that we were lucky that there wasn't any major accident after the Second World War. But this is not because of the presence of the United Nations. It is the counterbalance between the two superpowers. So I think we need some kind of balancing mechanism which would work naturally rather than an artificial organization like the United Nations. I don't think the United Nations is powerful enough and we got enough mechanisms.

Hu Jiaqi:

I agree on both points.

Here's how I think about it:

Unifying humanity is unprecedented and very difficult, but two ideas support my proposal.

First, the expansion of political entities is a major trend. In Chinese history, the Xia Dynasty was tiny, then Shang, Zhou, and Qin expanded.

World history shows the same—Mongol conquests, the British Empire. Overall, political entities tend to grow.

Second, the only hardware obstacle to expansion has been poor transportation and communication, which made large areas ungovernable. That's no longer the case.

Today's technology and communications have made the world a global village. If people form a consensus, the hardware conditions for unification are already met.

In the past, expansion relied on conquest. But there are peaceful examples:

Take the reunification of East and West Germany, for example. It was a peaceful process; no one forced them to do it.

Then look at the United States after its War of Independence. Back then, the 13 states were independent, and the country was a confederation, not the federal system we see today. Each of those 13 states was sovereign. Their decision to unite was entirely voluntary.

Now we must explain this clearly. Some friends joke that only an alien invasion would unify humanity. He made a joke out of my point about unity. But I said, 'You know what, you're absolutely right. When aliens invade Earth, humanity will unite. Why? Because the threat becomes obvious. People will immediately see that if we stay divided while being invaded, we're finished.

But we face the same danger from science and technology. If we don't unify under a global authority to firmly control technology, we will go extinct. This truth requires deep thinking to grasp. It's not that simple. It takes deep thinking to truly grasp this truth.

So my approach is to launch an awakening and enlightenment movement—to help people rethink science and technology, separate the good from the bad, spread the benefits globally to ensure prosperity, and firmly restrict the harmful aspects.

I founded the Humanitas Ark and have devoted my life to this awakening. Sharing these ideas with you, Professor, today is part of that effort. If everyone understands this, human unification is possible—just like German reunification or the union of the 13 U.S. states. It's possible. That's my point.

VI. Discussion on the Direction and Control of Technological Development

Konstantin:

I think we need to find the, uh , internal mechanisms how we can protect ourselves against the technology running away. So we were reasonably successful with the nuclear weapon. But, but they only have been what eighty years now. I'm sure that there will be a major incident somewhere in the future. But we definitely need to think, at least to think about the possible mechanisms how we can protect ourselves. So is this book translated to to English? I would I'll try to buy it on the amazon.

Hu Jiaqi

I'll give you a copy.

Konstantin:

I will definitely read it. And I will I will prepare better for our next meeting.

Hu Jiaqi:

Good. I also hope professor will share our views with your friends—let them know that in China, there is someone who has specialized in this issue for nearly half a century.

Konstantin:

I'll definitely read it through. And I'll definitely have another discussion after I've read it. And after that, I'll try to spread this message across my friends.

Hu Jiaqi:

Thank you. If you still have time, I'd like to say a few words about nuclear weapons.

Konstantin:

My view is that science cannot be stopped. So we cannot stop exploring. Unfortunately, sometimes technology which is created from science is dangerous. I don't think that can be avoided. But we can definitely have to limit the harm and have to control it after it has been created. I don't really know what is the best mechanism.

Hu Jiaqi:

Under today's social structure, science and technology cannot stop—because nations are independent and competitive. No country can halt its own technological progress.

But when humanity is on the brink of extinction, it's like someone drowning. Even without any hope, we will grasp at even a single straw.

Take Qin Shi Huang, for example. He was the first to unify China. Before that, China had been engulfed in warfare for hundreds of years and had never been truly united. No one believed unification was possible, yet he made it happen. When the tide of unification comes, nothing can stand in its way.

One or two nuclear bombs can't wipe out humanity. Neither can 200. Weapons that can cause extinction with just one or two uses—those are the real dangers.

The world is vast, with over 8 billion people today. Although the number of bullets and shells far exceeds 8 billion, they cannot wipe out humanity. This is because any extinction-level event would also destroy the user themselves. True extinction methods only need to be used once or twice—or require a very small amount—to annihilate mankind; only then can they be considered true extinction means.

The reason nuclear weapons can be controlled is threefold: First, only nations can acquire them. Second, it would likely take countless uses to actually wipe out humanity. And third, only the power of a state can possess them.

However, with today's fears about AI and the unchecked replication of nanotechnology, it is entirely possible that a single individual could operate them to exterminate the entire human race. And a single person is incredibly difficult to control. That is my view.

Therefore, I believe humanity has reached a critical juncture. It is already difficult to take measures to control this issue today, but if we don't act now, humanity will have even less of a future.

Konstantin:

Well, clearly it's not on the science side. So the science cannot be stopped. So it has to be done, retrospectively after the technology is already created. But I don't have a good solution here. So, I don't think the United Nations would would help. Up to now, the deterrent , was the , was the use of the... They was the keeper of the balance. I don't know what kind of protection can we take against AI. I think there there is some threat. But I don't know what kind of protection can can we take. I will be very curious to learn from your book.

Hu Jiaqi:

We are doing this work knowing it may be impossible. But we have no choice—if we fail, our species is gone.

Even if scientists' estimates vary, they agree: humanity doesn't have much time. If we're willing to go extinct, we can

ignore it. But a dying person grasps at any straw. That's what I've been doing for nearly half a century.

VII. Introduction to the Future Development Plan for Humanitas Ark

Hu Jiaqi:

In about a year, we plan to establish a Global Center for the Study of Technological Risks in the U.S. I'd like to hear your advice.

Konstantin:

I think there are several organizations already for the protection of humanity against the nuclear weapon. So I am part of those organizations. So I can show you how they function now. And then you can probably got some ideas. How it can be created. I can write some recommendations and send you the some ideas.

Hu Jiaqi:

Thank you very much. Currently, two points are not widely accepted globally:

First, the definitive argument: The continued development of science and technology will inevitably lead to the extinction of humanity in the very near future—at most two or three hundred years from now, and possibly within just one hundred years. This definitive argument is currently not widely accepted around the world.

Second, what is also not widely accepted: to solve this problem, humanity must achieve great unification—using the power of a world regime to firmly restrict the development of science and technology. This point is not widely accepted either. Those are the two points.

Konstantin:

I agree that there is a threat. And clearly even with nuclear power. There is quite a significant threat already. And I agree that dialogue between nations is something which will help us to avoid this disaster. I really like it here. So how to limit the development of science and technology. And I don't think it is possible in this society. So I completely agree. It's very difficult. It's impossible to limit the development of science and technology. That's a very good thought, yeah.

Hu Jiaqi:

For our Global Center for Technological Safety Research, we plan to invite top scientists—including those who disagree with us.

For example, in our very first conversation, Nobel laureate Michael Levitt argued that artificial intelligence might be capable of controlling its own risks. I remain reserved on this issue, but as long as he's thinking about technology and humanity's future, I respect that.

Konstantin:

Sure , absolutely. I don't really know, I don't have an opinion here to whether AI is a is a threat or not. But there is such a such a possibility. So , but I agree that you need people with different perspective, and with different views to to communicate together.

Hu Jiaqi:

The great Chinese philosopher Confucius once said, "Gentlemen seek harmony but not uniformity." As long as we share the same concerns and reflections for the future of humanity, even with differing views, we can be good friends.

Konstantin:

Yeah , that's a good point. So thank you so much for the book.